Learning to Segment Object Proposals via Recursive Neural Networks

Tianshui Chen, Liang Lin, Xian Wu, and Xiaonan Luo

arXiv:1612.01057v2 [cs.CV] 6 Dec 2016

***Abstract*—To avoid the exhaustive search over locations and scales, current state-of-the-art object detection systems usually involve a crucial component generating a batch of candidate object proposals from images. In this paper, we present a simple yet effective approach for segmenting object proposals via a deep architecture of recursive neural networks (RNNs), which hierarchically groups regions for detecting object candidates over scales. Unlike traditional methods that mainly adopt fixed similarity measures for merging regions or finding object proposals, our approach adaptively learns the region merging similarity and the objectness measure during the process of hierarchical region grouping. Specifically, guided by a structured loss, the RNN model jointly optimizes the cross-region similarity metric with the region merging process as well as the objectness prediction. During inference of the object proposal generation, we introduce randomness into the greedy search to cope with the ambiguity of grouping regions. Extensive experiments on standard benchmarks, e.g., PASCAL VOC and ImageNet, suggest that our approach is capable of producing object proposals with high recall while well preserving the object boundaries and outperforms other existing methods in both accuracy and efficiency.**

***Index Terms*—Object proposal generation, Object segmentation, Region grouping, Recursive neural networks, Deep learning.**

I. Introduction

OBJECT proposal generation, which targets on providing an ideal number of candidate regions where objects are likely to occur, benefits a wide range of applications such as generic object detection [1], [2] and object discovery [3], [4]. Usually, a good object proposal method is desired to be capable of recalling all existing objects over scales and locations while preserving their boundaries.

The challenges of object proposal lie in the presence of severe occlusion, variations in object shapes, and the lack of category information. Most of the current methods [5]–[7] tackle these difficulties through bottom-up region grouping or segmentation. Those methods mainly involve two crucial components, i.e., cross-region similarity metric and region merging algorithm. The similarity metric is utilized to measure whether two adjacent regions should be merged, and the merging algorithm performs the inference process that groups pairs of regions into super-regions and finally generates object proposals. Thus, object proposal generation methods based on region grouping basically follow the pipeline: they assign a higher similarity score to the adjacent regions if it is confident that the regions belong to the same class, and recursively merge the adjacent regions with highest score. Despite of acknowledged successes, these approaches usually require elaborative tuning or setting (e.g., manually designed cross-region similarity metric), limiting their performance in complex environments.

The authors are with the School of Data and Computer Science, Sun Yat-sen University, Guangzhou 510006, China. E-mail: linliang@ieee.org.

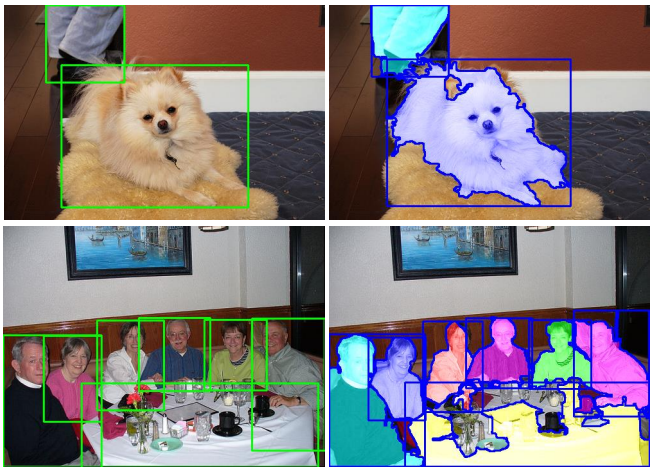

Fig. 1. Some object proposals (indicated by the blue boxes) generated by our approach. Our results match well with the ground-truth (indicated by the green boxes), and also preserve the object boundaries (indicated by the blue silhouettes inside the boxes).

In this work, we develop a novel hierarchical region grouping approach for generating and segmenting object proposals based on a recursive neural network (RNN). In our RNN architecture, we incorporate the cross-region similarity metric learning into bottom-up region merging process for end-to-end training. In particular, we define a structured loss that penalizes the incorrect merging candidates by measuring the similarity of adjacent regions and the objectness. In this way, the RNN explicitly optimizes the cross-region similarity learning and objectness prediction within the recursive iterations. Interestingly, the forward process of RNN finely accords with the traditional bottom-up region grouping pipeline, leading to a very natural embedding of the two crucial components (i.e., cross-region similarity metric and merging algorithm). Moreover, the objectness score is also learned with the RNN training, bringing the benefit of fast rejecting the false positive samples.

Obviously, the greedy merging algorithms, that recursively merge two regions with highest merging scores, can be applied for inference with the RNN model [8]. However, the performance of greedy methods depends heavily on the accuracy of merging scores, since greedy merging is generally

sensitive to noise or local minima. In the task of object proposal generation, once a segment of an object is incorrectly merged with the background or other objects, this object has little possibility to be recalled. In addition, we experimentally found that greedy merging leads to incorrect object proposals easily, especially when one segment of an object has similar appearance with background or other surrounding objects. To alleviate this issue, we propose a randomized merging algorithm that introduces randomness in the recursive inference procedure. Instead of merging a pair of neighbouring regions with highest similarity score, we search for k pairs with top k highest similarities, and then randomly pick one pair according to a distribution constructed by their scores. This helps to recall more incorrectly merged objects. Figure 1 shows some examples of object proposals generated by our approach.

The key contributions of this work a deep architecture of recursive neural networks for generating object proposals and preserving their boundaries. This framework jointly optimizes the cross-region similarity and objectness measure together with the hierarchical region grouping process, which is original in literature of object segmentation and detection. Moreover, we design a randomized region merging algorithm with the recursive neural network learning, which introduces randomness to handle the inherent ambiguities of composing regions into candidate objects and thus causes a notable gain in object recall rate. Extensive experimental evaluation and analysis on standard benchmarks (e.g., PASCAL VOC and ImageNet) are provided, demonstrating that our method achieves superior performances over existing approaches in both accuracy and efficiency.

The remainder of the paper is organized as follows. Section II presents a review of the related works. We then introduce our approach and optimization algorithm in detail in Section III and Section IV, respectively. Experimental results, comparisons and analysis are exhibited in Section V. Section VI concludes the paper.

II. Related Work

Many efforts have been dedicated to object proposal generation. Here we roughly divide existing methods into two categories: top-down window-based scoring and bottom-up region grouping, according to their computation process.

A. Window-based Scoring

This category of methods [9]–[13] attempt to distinguish object proposals directly from the surrounding background through assigning an objectness score to each candidate subwindow. The objectness measures are usually defined in diverse ways, and object proposals generated by sliding windows are then ranked and thresholded by their objectness scores. As a pioneer work, Alexe et al. [12] employed saliency cue to measure the objectness of a given window, which was further improved by [14] with learning methods and more complicated features. However, these methods may suffer from expensive computational cost, since they require to search over all locations and scales in images. Recently, to address this problem, BING [11] and Edge Box [10] exploited very simple features such as gradient and contour information to score the windows, and achieved very high computational efficiency. Alternatively, Ren et al. [9] proposed a deep learning method based on fully convolutional networks (FCNs) [15] to score windows over scales and locations efficiently. Nonetheless, this method may not locate object accurately, since experimental results show that the recall rate deteriorates as the Intersection over Union (IoU) threshold increases.

B. Region Grouping

This branch of researches [5], [6], [16]–[20] cast the object proposal generation as a process of hierarchical region segmentation or partition. Starting from an initial oversegmentation, these methods usually adopt a cross-region similarity / distance metric that works together with region merging algorithms. As a representative example of these methods, Uijlings et al. [5] leveraged four types of low-level features (e.g., color, texture etc.) for similarity computing and generated object proposals via hierarchical greedy grouping. Using similar features with [5], Manen et al. [6] learned the merging probabilities and introduced a randomized prim algorithm for region grouping. Arbeláez et al. [18] constructed hierarchical segmentations and explored the combinatorial space to combine multi-scale regions into proposals. Rantalankila et al. [19] integrated local region merging and global graph-cut to generate proposals. Due to their high localization accuracy, they are adopted in many state-of-the-art object detection [1], [2] and object discovery [4] algorithms. However, these mentioned methods mainly adopt fixed similarity measures for merging regions or finding object proposals, leading to suboptimal performances when handling complex cases. In contrast, our approach adaptively learns the region merging similarity and the objectness measure during the process of hierarchical region grouping. Moreover, our method also introduces randomness into the bottom-up searching of region composition and yields significant improvement over existing methods.

III. Framework of Segmenting Object Proposals

In this section, we introduce our approach in detail. The input image is first over-segmented into N regions with the efficient graph-based method [21]. The Fast R-CNN [22] is used to extract local features for each region. We then design a recursive neural network to group regions and simultaneously predict the associated objectness scores for corresponding proposals. Furthermore, we propose a randomized merging algorithm, which introduces randomness into recursive inference procedure to cope with the inherent ambiguities in the process of merging regions. Figure 2 gives an illustration of our proposed framework.

A. Local Feature Extraction

We utilize the Fast RCNN [22] architecture to extract local features for each region. The architecture consists of 16 convolutional layers, the same as VGG16-net [23], followed by the region of interests (ROI) pooling layer. Specifically,

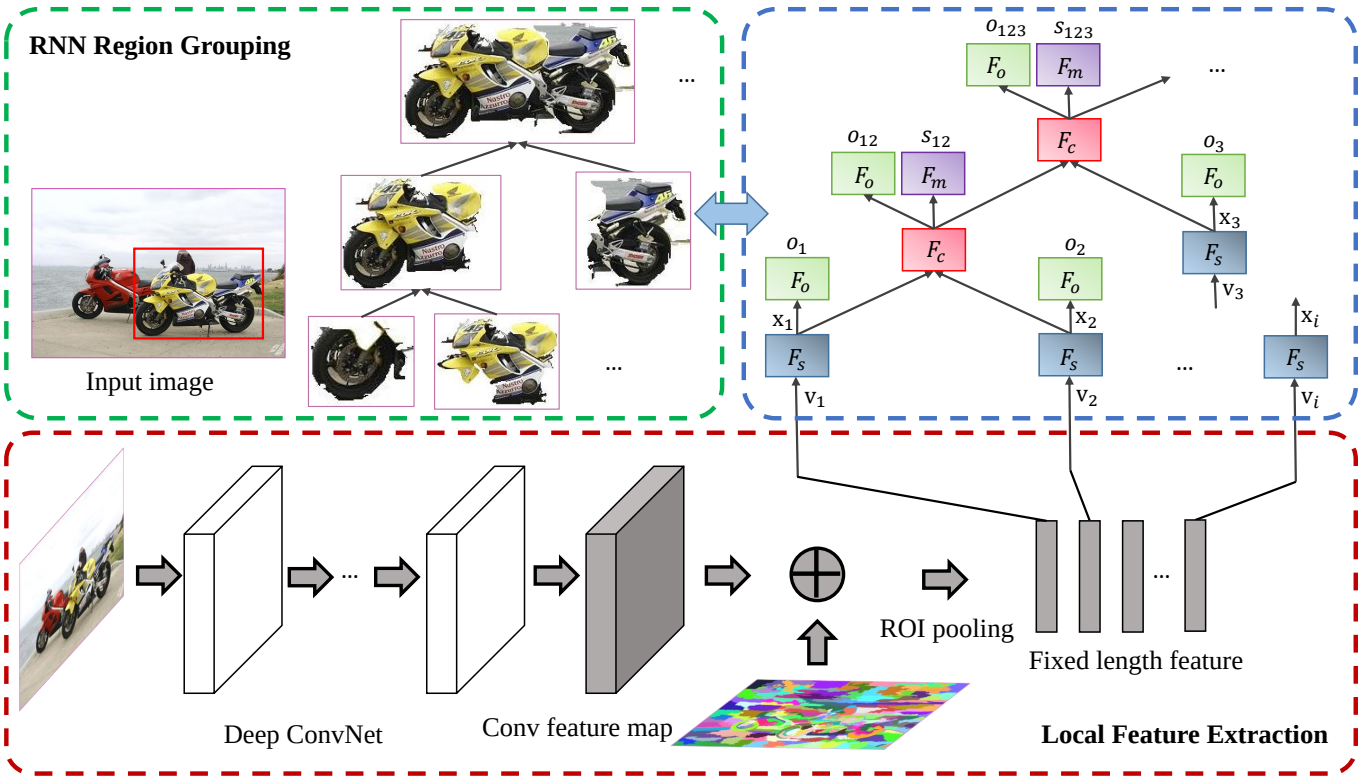

Fig. 2. An overview of our proposed object proposal segmentation framework. The bottom shows local feature extraction, and the top illustrates bottom-up recursive region grouping process. The four modules, F_s, F_c, F_m and F_o, work cooperatively to group regions for generating object proposals.

given an input image, our approach first over-segments it into N regions with the efficient graph-based method [21] and obtains the box for each region that tightly bounds this region. The image is resized thus that the short side is 600, remaining the aspect ratio unchanged. The sixteen convolutional layers take the resized image as input, and produce a pooling of corresponding size feature maps. The ROI pooling layer subsequently extracts a fixed length feature vector for each region.

B. Recursive Neural Network

We first present some notations that would be used throughout this article. Let $\mathbf{v}_i$ denote the local features of the i-th region, and $\mathbf{x}_i$ denote the corresponding semantic features. $\sigma(\cdot)$ denotes the rectified linear unit (ReLU), where $\sigma(x) = \max(0, x)$.

The core of this framework is the RNN, which aims to group the regions and simultaneously predict the objectness scores for corresponding proposals in a recursive manner. The RNN architecture is depicted in Figure 3. The RNN comprises four modules, i.e., semantic mapper, feature combiner, merging scorer and objectness scorer. Semantic mapper transforms the local features to semantic space which can be further propagated to their parent nodes. Feature combiner computes the joint semantic representations of all neighbouring child nodes. Given joint semantic representations, merging scorer calculates the score indicating the confidence that two nodes should be merged. Feature combiner merges the neighbouring nodes according to merging scores, and obtains a hierarchical tree structural segmentations, each of which corresponds to a candidate of proposal. Objectness scorer computes a score which estimates the likelihood of the candidate containing an object. These four modules work cooperatively for proposal segmentation, as illustrated in Figure 2. We describe these four modules in the following.

1) Semantic Mapper: Semantic mapper F_s is a simple feed-forward operator to map the local features into the semantic space in which the combiner operates on. It can be expressed as,

$$\mathbf{x}_i = F_s(\mathbf{v}_i; \theta_s) = \sigma(W_s \mathbf{v}_i + \mathbf{b}_s), \quad (1)$$

where θ_s are the learned parameters. The semantic mapper captures the region semantic representation, and propagates it to its parent regions through the tree hierarchical structure. To better balance the computational efficiency and accuracy, we empirically set the dimensionality of local features $\mathbf{v}_i$ as 18,432 ($6 \times 6 \times 512$), and that of semantic features $\mathbf{x}_i$ as 256. Hence, the semantic mapper is a one-layer fully-connected network, with 18,432 input and 256 output neurons, followed by the rectified linear unit.

2) Feature combiner: Feature combiner F_c recursively takes the semantic features of two child nodes as input, and maps them to the semantic features of the parent node, formulated as,

$$\mathbf{x}_{i,j} = F_c([\mathbf{x}_i, \mathbf{x}_j]; \theta_c) = \sigma(W_c[\mathbf{x}_i, \mathbf{x}_j] + \mathbf{b}_c), \quad (2)$$

Feature combiner aims to aggregate the semantic information of the two child nodes and obtains the semantic representation

of the merged node. It takes semantic features of the original regions as leaf nodes, and recursively aggregates them to the root node in a bottom-up manner. In order to ensure the recursive procedure can be applied, the dimensionality of parent node features is set the same as that of child node features. Thus, the architecture of the feature combiner is identical to that of the semantic mapper, except that it has 512 (2×256) input neurons.

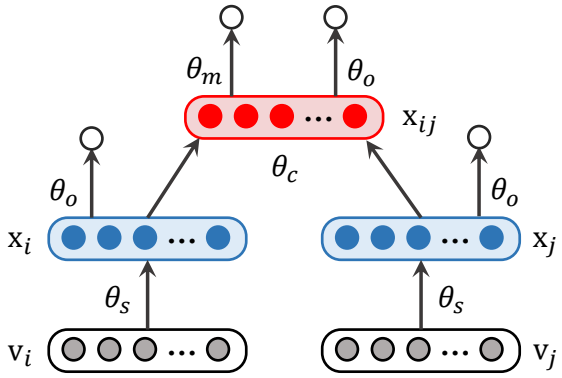

Fig. 3. Illustration of the recursive neural network in our proposed framework. This network computes the scores for merging decision and objectness scores of all regions.

3) Merging scorer: Given joint semantic features of two neighbouring nodes, merging scorer F_m computes a score that indicates the confidence that whether two nodes should be merged, expressed as

$$s_{i,j} = F_m\left(\mathbf{x}_{i,j}; \theta_m\right) = W_m \mathbf{x}_{i,j}, \tag{3}$$

The scores determine the pair that should be merged first in both learning and inference stages. It consists of one simple fully connected layer which takes 256 dimensionality combined features as input and produces one score.

4) Objectness scorer: Each node of the tree is related to the semantic information of the corresponding region, i.e., the semantic features. Objectness scorer F_o directly predicts objectness scores in semantic feature space.

$$o_i = F_o\left(\mathbf{x}_i; \theta_o\right) = \phi\left(W_{o_1}\sigma\left(W_{o_0}\mathbf{x}_i + \mathbf{b}_{o_0}\right)\right), \tag{4}$$

where $\phi(\cdot)$ is the softmax operation. Our approach rejects candidate proposals that have low scores without compromising the recall rate. We experimentally found that one fully connected layer (merely consisting of 512 parameters) is so simple that it can not well fit thousands of proposals. Thus, we utilize two stacked fully connected layers to implement the objectness scorer, in which the first one is 256 to 256, followed by the rectified linear unit, and the second one is 256 to 2, followed by a softmax layer for objectness prediction.

C. Randomized Merging Algorithm

As discussed above, greedy merging groups the neighbouring regions with the highest similarity score for each iteration. Once a segment of an object mistakenly merges with a neighboring segment that belongs to surrounding objects or background, this object would have little chance to be found. Figure 4 presents an example as an illustration. Given an image with a brown cat and a black-white one, the brown cat is successfully detected using the greedy merging processing, as it is distinguishable from the background (red bounding box in Figure 4). However, the white segment of the other cat incorrectly merges with a piece of background as they have more similar appearance (red circle in Figure 4). In this case, the subsequent merging process misses this cat inevitably. We propose the randomized merging algorithm to cope with this problem. Instead of merging the neighbouring regions with the highest similarity score for each iteration, our approach selects one pair to merge among the top k highest pairs according to a distribution constructed based on their scores. The randomized merging process can be repeated for several times to increase the diversity of the generated proposals. This helps to recall more incorrectly merged objects, as explained in Section V-E.

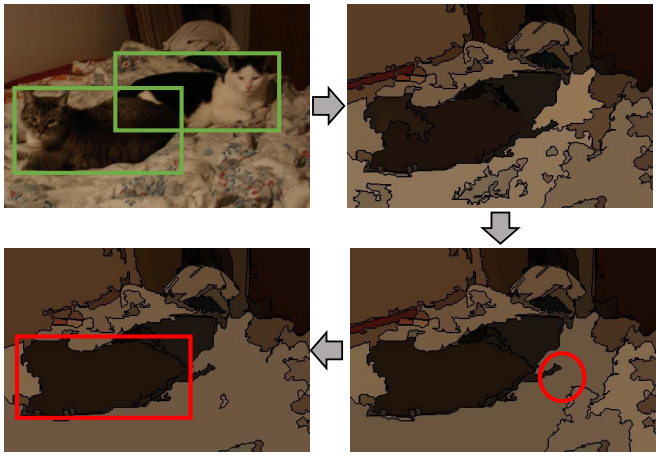

Fig. 4. An example of incorrect merging using the greedy merging algorithm. Top left: Input image; top right: over-segmentation; bottom right: incorrect merging; bottom left: merging result. The black-white cat is lost because its white part incorrectly merges with the background.

The randomized merging algorithm works as follows. Starting from the semantic features $\{\mathbf{x}_i\}_{i=1}^{N_{seg}}$ and over-segmented regions $\mathcal{R} = \{r_i\}_{i=1}^{N_{seg}}$, our approach first computes the merging scores of all neighbouring regions using the feature combiner and merging scorer. Our approach then re-ranks the merging scores to obtain the k neighbouring regions $\{(r_{i_t}, r_{j_t})\}_{t=1}^{k}$ with top-k highest scores $\{s_{i_t,j_t}\}_{t=1}^{k}$, and further constructs a multinomial distribution according to the k merging scores, expressed as

$$(i_t, j_t) \sim Mult\left(\rho\right), \tag{5}$$

where

$$\rho_{i_t,j_t} = \exp\left(s_{i_t,j_t}\right) / \sum_{t=1}^{k} \exp\left(s_{i_t,j_t}\right). \tag{6}$$

Our approach randomly picks one pair of regions $\left(r_{i_{t'}}, r_{j_{t'}}\right)$ according to the distribution $Mult\left(\rho\right)$, merges these two regions together, and then computes new merging scores between the resulting region and its neighbours. The process is repeated until the whole image becomes one region. The general process is detailed in Algorithm 1.

Algorithm 1 Randomized merging algorithm

Input: Initial region set $\mathcal{R} = \{r_i\}_{i=1}^{N_{seg}}$
Output: Set of object proposal $\mathcal{P}$

Initialize merging score set $\mathcal{S} = \emptyset$
for all neighbouring region pair (r_i, r_j) **do**
  Calculate merging score $s_{i,j}$
  $\mathcal{S} = \mathcal{S} \cup s_{i,j}$
end for
while $\mathcal{S} \neq \emptyset$ **do**
  Get the k highest merging scores $\{s_{i_t,j_t}\}_{t=1}^{k}$
  Construct multinomial distribution $Mult(\rho)$
  Select randomly t'-th pair according to $Mult(\rho)$
  Merge corresponding regions $r_{t'} = r_{i_{t'}} \cup r_{j_{t'}}$
  Remove scores regarding $r_{i_{t'}}$: $\mathcal{S} = \mathcal{S} \setminus s_{i_{t'},*}$
  Remove scores regarding $r_{j_{t'}}$: $\mathcal{S} = \mathcal{S} \setminus s_{j_{t'},*}$
  Compute merging score set $\mathcal{S}_{t'}$ between $r_{t'}$ and its neighbours
  Update merging score set $\mathcal{S} = \mathcal{S} \cup \mathcal{S}_{t'}$
  Update region set $\mathcal{R} = \mathcal{R} \cup r_{t'}$
end while
Extract object proposals $\mathcal{P}$ from all regions in $\mathcal{R}$

IV. Optimization

Suppose that we are given the training set $\mathcal{X} = \{(I_i, c_i, b_i) \,|i = 1, 2, ..., N\}$, where N is the number of training samples; I_i is the i-th input sample, including the local features of all regions and the adjacency matrix (as shown in Figure 5(a) and (b)); c_i and b_i are the corresponding class labels of regions and ground truth object bounding boxes, respectively. Our model is jointly trained with two objectives: 1) the merging loss $\mathcal{L}_m$ penalizes incorrect region grouping in the hierarchical tree structure; and 2) the objectness loss $\mathcal{L}_o$ helps to learn the objectness scorer. Therefore, we define the structured loss as

$$\mathcal{L} = \mathcal{L}_m + \lambda \mathcal{L}_o + \frac{\eta}{2}||\theta||_2^2, \tag{7}$$

where $\theta = \{\theta_s, \theta_c, \theta_m, \theta_o\}$ are the set of parameters to learn and $||\theta||_2^2$ is the regularization term. λ and η are two balance parameters.

A. Merging Loss

Given an input sample I, its bottom-up merging process can be presented as a binary tree $t \in \mathcal{T}(I)$, where $\mathcal{T}(I)$ is the set of all possible binary trees constructed from input I. In the learning stage, the class labels of all the segmented regions are available. We further define $\mathcal{T}(I, c)$ as the set of all possible correct trees. Here, a tree is regarded as correct if any region merges with the one belonging to the same class before other regions from different classes. Figure 5 presents some examples of generating correct and incorrect trees from an image.

Inspired by [8], [24], we define a margin loss function $\triangle L : \mathcal{I} \times \mathcal{C} \times \mathcal{T} \rightarrow \mathbb{R}^+$, where $\triangle L(I, c, t)$ measures the penalty of the construction of a parsing tree t for input I with label c. In the context of recursive merging process, the loss

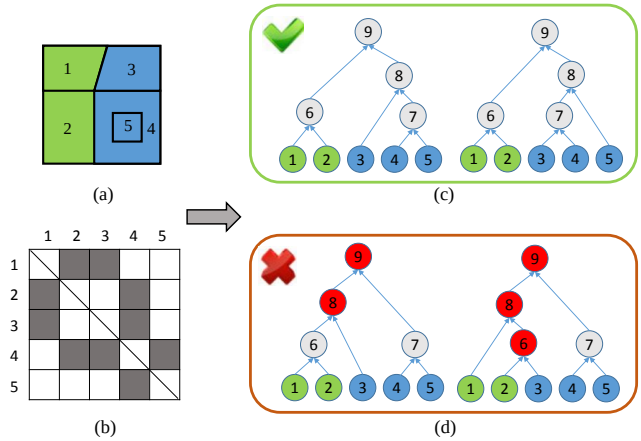

Fig. 5. Some examples of generating correct and incorrect trees. (a) Input image, green and blue indicate differently labelled regions. (b) Adjacent matrix of image regions; (c) correct trees; (d) incorrect trees.

increases when a segment merges with the one from different class before those with the same class label. We denote $N(y)$ as the set of non-terminal nodes of tree t, and $subtree(d)$ as a subtree underneath the non-terminal node for each $d \in N(y)$. Naturally, we formulate the loss by penalizing the incorrect subtrees

$$\triangle L(I, c, t) = \sum_{d \in N(y)} \mathbf{1}\{subtree(d) \notin \mathcal{T}(I, c)\}, \tag{8}$$

where $\mathbf{1}\{\cdot\}$ is an indicator function whose value is 1 when the expression is true and 0 otherwise. Figure 5(d) illustrates two examples of incorrect trees, in which the margin losses are 2 and 3, respectively.

Our goal is to learn a function $f_\theta(\cdot)$ with small expected loss on the unseen inputs. Similar to [8], [24], we consider the following forms

$$f_\theta(I) = \arg\max_{t \in \mathcal{T}(I)} \{s(RNN(\theta, I, t))\}, \tag{9}$$

where $s(\cdot)$ predicts the score for a tree by summing up merging scores of all the merged neighbouring pairs. In the optimization procedure, we aim to learn a score function that assigns higher scores to correct trees than incorrect ones. Given the parameters θ, we first define the margin between the correct tree t_i and another tree t for I_i,

$$s(RNN(\theta, I_i, t_i)) - s(RNN(\theta, I_i, t)). \tag{10}$$

Intuitively, the margin will be enlarged as the margin loss function $\triangle L(I, c, t)$ increases, expressed as

$$s(RNN(\theta, I_i, t_i)) - s(RNN(\theta, I_i, t)) \geq \kappa \triangle L(I, c, t), \tag{11}$$

where κ is a parameter. So the merging loss can be defined as

$$\mathcal{L}_m = \sum_{i=1}^{N} \mathcal{L}_m^{(i)}, \tag{12}$$

where

$$\mathcal{L}_m^{(i)} = \max_{t \in \mathcal{T}(I_i)} \left\{ s\left(RNN\left(\theta, I_i, t\right)\right) + \kappa \triangle L\left(I_i, c_i, t\right) \right\} - \max_{t_i \in \mathcal{T}(I_i, c_i)} \left\{ s\left(RNN\left(\theta, I_i, t_i\right)\right) \right\}. \tag{13}$$

Optimizing the merging loss can maximize the correct trees' scores while minimizing the scores of the highest scoring but incorrect trees.

B. Objectness Loss

One of the main advantages of our approach is that it can simultaneously predicts an objectness score for each proposal candidate, which can be used for proposal ranking and rejecting the ones with low scores. We simply employ a softmax classifier with the semantic features of each node. We generate positive and negative samples from all of the regions as follows. Given a region, we first calculate the IoU scores between the box that tightly bounds this region with each ground truth bounding box. If the maximum IoU is larger than 0.5, this region is considered as positive; and if the maximum IoU is smaller than 0.2, it is used as a negative sample. We simply ignore other regions since they may not provide discriminative information. For the i-th useful region, the loss function can be defined as

$$\mathcal{L}_o^{(i)} = -\sum_{l=0}^{1} \mathbf{1}\left\{l_i = l\right\} \log\left(p_{i,l}\right), \tag{14}$$

where $p_{i,l}$ is the score that corresponds to the likelihood that the region belongs to label l. Hence

$$\mathcal{L}_o = \sum_{i=1}^{N_u} \mathcal{L}_o^{(i)}, \tag{15}$$

where N_u is the number of useful regions.

The model is jointly trained by the stochastic gradient descent (SGD) with momentum [25].

V. Experiment

In this section, we present the extensive experimental results to compare with state-of-the-art methods, demonstrating the superiority of the proposed methods, and analyze the benefit of introducing the randomized merging algorithm for object proposals generation.

A. Experimental Setting

1) Datasets: We first conduct the experiments on the PASCAL VOC2007 dataset [26], which consists of 9,963 images from 20 categories of objects. The model is trained using 422 images of the PASCAL VOC2007's segmentation set. We compare the performance of our approach with those of state-of-the-art methods, and evaluate the contribution of randomized merging algorithm using the 4,952 test images that contain 14,976 objects, including the "difficult" ones. To better demonstrate the effectiveness of the proposed method, we also conducted experiments on the PASCAL VOC 2012 validation set, which contains 15,787 objects in 5,823 images. As our model is trained with 20 object categories on PASCAL VOC, we further investigate the generalization ability of our method to unseen object categories on ImageNet 2015 validation dataset [27], which contains about 20,000 images of 200 categories, without re-training the model using the training samples from ImageNet.

2) Evaluation Metrics: One of the primary metrics is the Intersection over Union (IoU) measure, where the IoU is defined as the intersection area of the proposal, and the ground truth bounding box divided by their union area. For a fixed number of proposals, the recall rate (the fraction of ground truth annotations covered by proposals) varies as the IoU threshold increases from 0.5 to 1, so that a recall-IoU curve can be obtained. Besides, the curves indicating the recall rate with reference to the number of ranked proposals, are also given, with IoU fixed as both 0.5 and 0.8, respectively. This is widely adopted by many proposal works [28], [29] for evaluation. We also compare the average recall (AR), defined as the average recall when IoU ranges from 0.5 to 1 [28], [30], since AR is considered to be strongly correlated with detection performance.

3) Implementation Details: Following [5], we adopte the efficient graph-based method [21] to produce initial over-segmentations with four parameter values (i.e., $k = 100, 150, 200, 250$), respectively. We implement the proposed model using Caffe open source library [31], and train it by stochastic gradient descent (SGD) with a batch size of 2, momentum of 0.9, weight decay of 0.0005. The learning rate is initialized as 10^{-5} and divided by 10 after 20 epochs. Note that it is indeed possible to perform joint training for the fast RCNN and the RNN. We do not train the model in this way, because only 422 images are provided for training, which easily leads to over-fitting. Therefore, we simply use the fast RCNN model pre-trained on the PASCAL VOC 2007 detection dataset for local feature extraction, and then train our RNN model alone. In the random merging algorithm, the parameter k is set as 5. We conducted the experiments on a single NVIDIA GeForce GTX TITAN-X. It takes about 0.2s for feature extraction, about 0.6s for generating hundreds of object proposals. To improve the quality of generated proposals, we can allow the randomized merging process to be repeated several times (e.g., 8) in each recursion step.

B. Comparison With State-of-the-Art Approaches

In this subsection, we compare our method with recent state-of-the-art methods, including BING [11], Randomized Prim (RP) [6], EdgeBox (EB) [10], Multiscale Combinatorial Grouping (MCG2015) [1] [32], Selective Search (SS) [5], Faster R-CNN (RPN) [9], Complexity Adaptive Distance Metric (CADM) [7], Multi-branch Hierarchical Segmentation (MHS)[2] [29], Geodesic Object Proposals (GOP) [3] [16], Learn

[1]MCG2015 is the improved version of original MCG and achieves much better performance.

[2]We only compare with MHS on the PASCAL VOC 2007 dataset, because only the results on this dataset are available.

[3]We only compare with GOP on the PASCAL VOC 2007 and ImageNet datasets, because only the results on these two datasets are available.

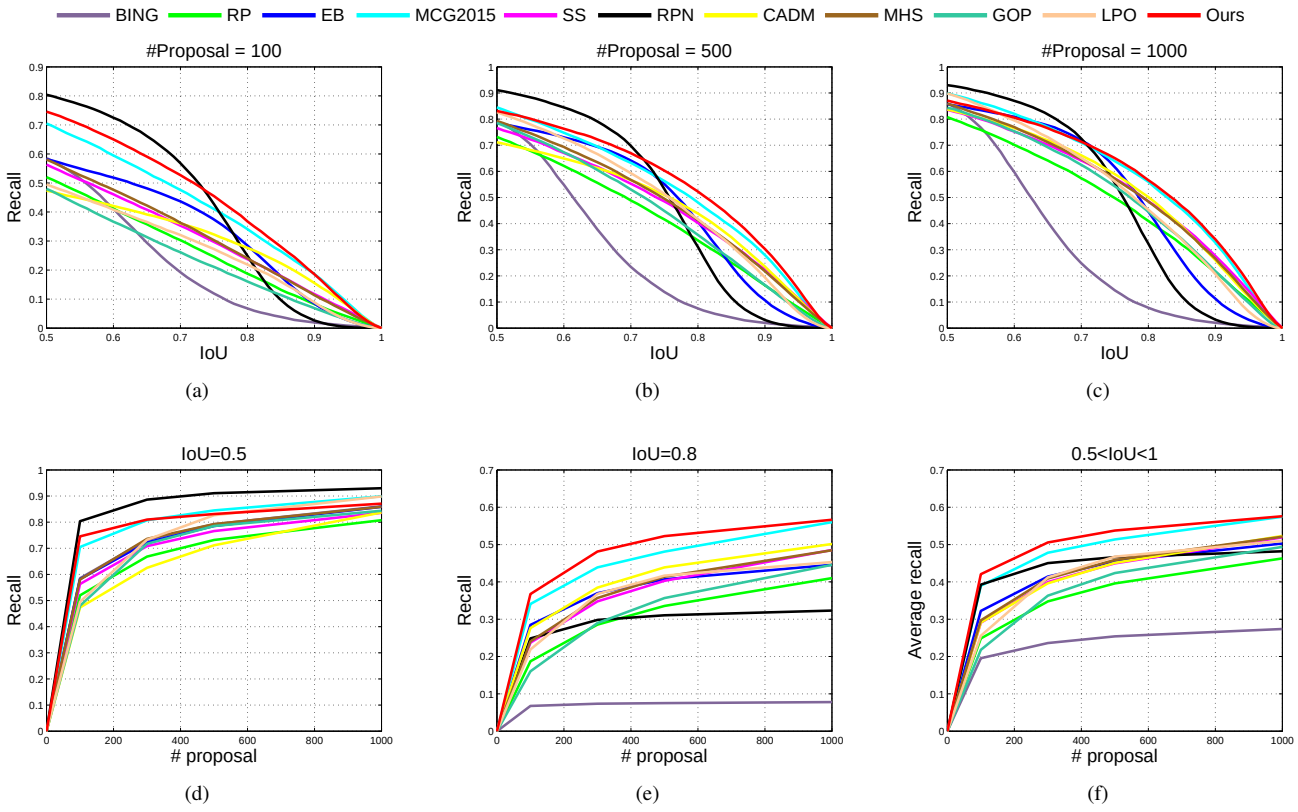

Fig. 6. Comparison of our proposed method and other state-of-the-art approaches on the PASCAL VOC 2007 test set. Best viewed in color.

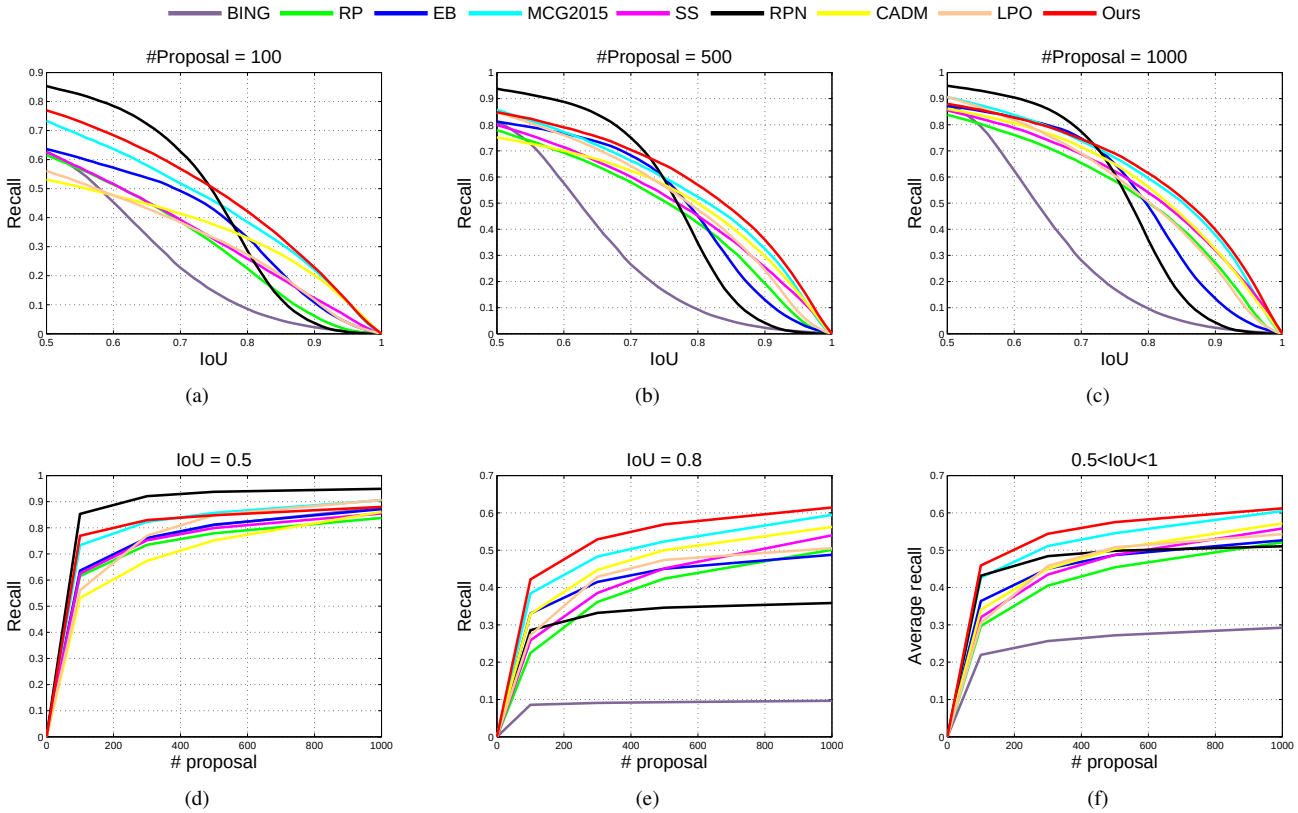

Fig. 7. Comparison of our proposed method and other state-of-the-art approaches on the PASCAL VOC 2012 validation set. Note that our model is trained on PASCAL VOC 2007, but still achieves the best performance against other competitors. Best viewed in color.

Fig. 8. Qualitative examples of our object proposals. Ground truth box are shown in green and red, with green indicating the object is found and red indicating the object is not found. The blue boxes are the object proposals with highest IoU to each ground truth box, and the blue silhouettes are the corresponding object contours. All the samples are taken from PASCAL VOC dataset.

to Propose Objects (LPO) [33]. In our experiments, we use Edgebox70 (optimal settings for an IoU threshold of 0.7) for EB, and default settings for others, in order to ensure the best overall performance for these methods. In addition, we follow [28] to control the number of candidates to a specific value for a fair comparison. Since BING, MCG2015, SS, CADM, RPN, EB, MHS and GOP provide sorted proposals, we select n proposals with top n highest scores for evaluation. However, RP and LPO does not provide the scores to rank the proposals, so we simply select the first n proposals in our experiments.

We first analyze the experimental results on the PASCAL VOC 2007 dataset, as depicted in Figure 6. It can be observed that window-scoring-based methods (e.g., EB and RPN) achieve competitive recall rates with a relatively low IoU threshold. This mainly benefits from the exhaustive search over locations and scales, and the accuracy of rejecting the non-objects by the window-scoring-based methods. However, their recall rates drop significantly when the IoU threshold increases. In contrast, region-grouping-based methods yield better performance as the IoU threshold increases. It is shown that MCG2015 performs best among region-grouping-based methods, but it is very time-consuming (over 30s per image) and may not practical especially for real-time object detection systems. It is noteworthy that our method runs 7$\times$ faster than

MCG2015, and meanwhile outperforms MCG2015 overall, particularly when the number of object proposals is strictly constrained (e.g., with 100 or 500 proposals).

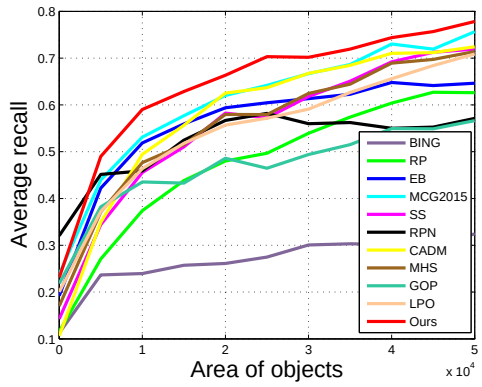

Fig. 9. Comparison of average recall (AR) with respect to different sizes of ground-truth objects on the PASCAL VOC 2007 test set. All of the AR rates are computed with top ranked 500 proposals per image. Best viewed in color.

Typically, an IoU threshold of 0.5 is used to measure whether the target object is detected successfully in object detection tasks. However, as suggested in recent works [10], [28], [30], the proposals with an IoU of 0.5 cannot fit the ground truth objects well, usually resulting in a failure of subsequent object detectors. This reveals the fact that the recall rate with an IoU threshold of 0.5 is weakly correlated with the real detection performance. Hence, we also present the curves of recall rate with respect to the number of proposals at a more strict IoU threshold of 0.8, shown in Figure 6(e), to demonstrate the superiority of our method. We believe that our method may be more suitable for object detection systems owing to better localization accuracy and efficiency. Besides, we compare the average recall (AR), considered to have a strong correlation with detection performance, as another important metric for evaluation. As shown in Figure 6(f), our method outperforms other state-of-the-art algorithms, which suggests that it is likely to achieve a better detection performance with the proposals generated by our method.

We also compare the performance on the PASCAL VOC 2012 validation set, as depicted in Figure 7. Note that RPN is trained with the data from both VOC 2007 and 2012 datasets, but RP and our method are learned on the VOC 2007 dataset without re-training here. Even though VOC 2012 is more challenging and larger in size, our method still achieves best performance over other state-of-the-art algorithms, again demonstrating the effectiveness of the proposed method. It is also shown that more improvement over other methods on the VOC 2012 is achieved than that on the VOC 2007.

We present some qualitative examples in Figure 8, including some random samples (top four rows) that contains two or more objects and some samples (the last row) that challenges our method. We find that, in most cases, our results match well with the ground-truth, and preserve the accurate object boundaries. The missed object are in part tiny ones, e.g., the distant and severely-occluded ones.

Figure 9 analyzes the AR with regard to the ground truth objects in different sizes. It is shown that our method performs slightly worse than RPN if we only consider small-sized objects whose areas are less than 5k pixels. Nevertheless, our method yields best performance over other competitors in general, especially for recalling objects in larger size. Other grouping-based methods such as SS and CADM show similar results. One possible reason is that grouping-based methods depends heavily on the over-segmentations. In this case, the boundaries of small-sized objects are generally difficult to be well preserved, if the segmentation is not accurate enough. But this problem will not have a significant impact on larger-sized objects. Therefore, region-grouping-based approaches usually exhibit desirable ability to recall objects in relatively large size, but may fail to recall small-sized ones.

C. Object Detection Performance

Since most state-of-the-art object detectors rely on object proposals as a first preprocessing step, it is essential to evaluate the final detection performance with proposals generated by different methods. In this subsection, we conduct experiments to analyze the quality of proposals for object detection tasks. To this end, we use the Fast R-CNN detection framework [22] using CaffeNet [34] as the benchmark. The detectors are trained using PASCAL VOC 2007 trainval set, and test using the test set for all the experiments here. We compare EB, SS, MCG2015 and RPN with the proposed method. For a fair comparison, we select top-1000 proposals for all of the methods in both training and testing stages. The mean average precision (mAP) and average precision (AP) for each of the 20 categories are shown in Table I. It can be seen that our proposed method achieves the best mAP of 58.6%, outperforming other state-of-the-art methods. This also verifies the effectiveness of our method for detection tasks.

D. Generalization to Unseen Categories

In addition, we conduct experiments on the ImageNet 2015 validation set to further evaluate the generalization ability to a wider scope of object categories. Note that all of the learning-based models (e.g., RP, RPN and ours) are trained on the PASCAL VOC dataset, and directly tested on the ImageNet 2015 validation set without re-training. The comparision of the experimental results are shown in Figure 10. It can be seen that our method has comparable performance with MCG2015, and surpasses other methods. No obvious deterioration in performance is observed on the ImageNet 2015 validation set, suggesting that our method does not exclusively fit the 20 specific categories of objects from the PASCAL VOC. In other words, our method is capable of capturing generic objectness information and generalizing to unseen categories. In addition, most state-of-the-art methods achieve similar performance to those on the PASCAL VOC, while RPN suffers from severe performance drop. One possible reason is that category information is exploited to learn class-specific detectors, which makes the RPN overfit 20 categories of objects from the PASCAL VOC.

Methods	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP
SS	**64.8**	67.9	52.9	45.3	20.6	69.9	65.6	70.7	30.0	62.7	59.8	62.3	73.6	64.9	54.8	24.8	49.3	60.2	71.7	55.2	56.4
EB	62.7	68.0	52.5	**45.8**	24.6	65.2	69.2	**70.8**	29.8	64.2	59.3	61.7	**74.9**	67.5	59.5	26.8	51.6	55.2	71.1	56.9	56.9
MCG2015	64.2	70.6	50.3	42.5	26.4	70.8	66.4	69.3	29.8	63.7	61.2	61.0	72.7	65.5	57.7	28.1	50.3	61.1	70.3	59.1	57.0
RPN	61.5	**71.2**	53.5	42.6	**29.0**	**72.3**	**72.5**	70.7	32.2	63.5	56.0	62.8	**74.9**	**68.2**	**62.4**	27.6	**53.9**	53.6	66.7	58.5	57.7
Ours	64.6	67.1	**56.1**	44.0	27.7	70.1	66.7	70.5	**35.9**	60.5	**61.8**	**64.0**	74.2	66.1	57.7	**30.2**	52.6	**62.7**	**73.2**	**65.7**	**58.6**

TABLE I

COMPARISON OF OBJECT DETECTION PERFORMANCE WITH PROPOSALS GENERATED BY DIFFERENT METHODS. ALL OF THE DETECTORS ARE LEARNED BY THE FAST RCNN ON PASCAL VOC 2007 TRAINVAL SET, AND TESTED WITH 20 CATEGORIES ON THE PASCAL VOC 2007 TEST SET.

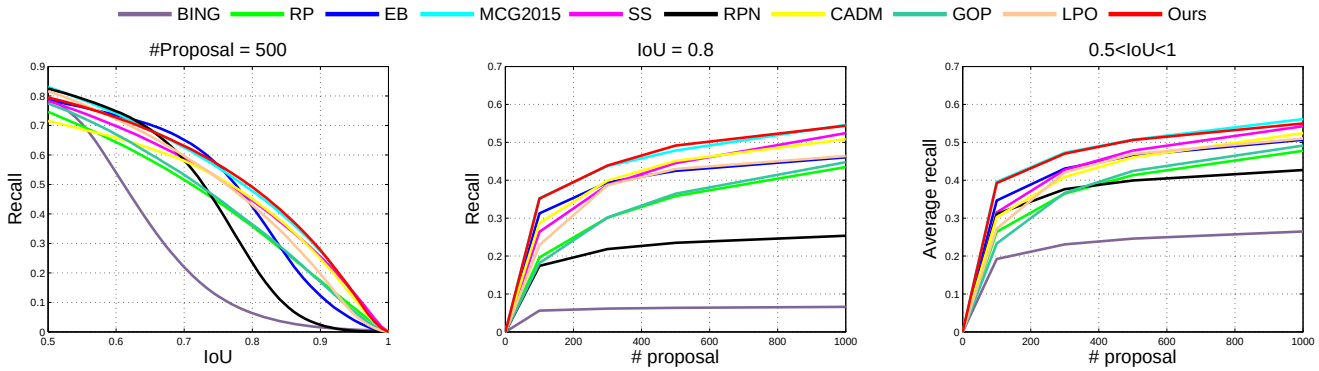

(a) (b) (c)

Fig. 10. Comparison of our proposed method and other state-of-the-art approaches on the ImageNet 2015 validation set. Best viewed in color.

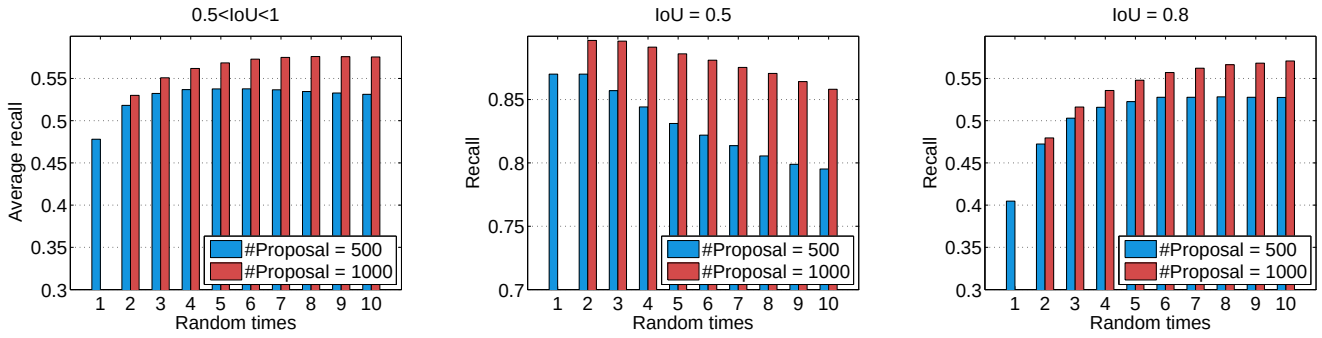

(a) (b) (c)

Fig. 11. Comparison of recall rates with different randomized times on the PASCAL VOC 2007 test set. We report the results of both top 500 and 1000 proposals for a fair comparison. Note that the number of proposals generated by one-time randomized merging is less than 1000, so we cannot provide this result here. Best viewed in color.

	R@0.5	R@0.8	AR
Greedy	0.872	0.423	0.489
Random	0.870	0.405	0.478

TABLE II

COMPARISON OF GREEDY MERGING AND RANDOMIZED MERGING ON THE PASCAL VOC 2007 TEST SET. WE REPORT THE RESULTS USING TOP RANKED 500 PROPOSALS. R@0.5 AND R@0.8 INDICATE THE RECALL RATES WITH AN IOU THRESHOLD OF 0.5 AND 0.8, RESPECTIVELY, AND AR IS THE AVERAGE RECALL.

E. Evaluation of Randomized Merging Algorithm

In this subsection, we evaluate the contribution of the proposed randomized merging algorithm. We compare the performance of conventional greedy and the proposed randomized merging algorithms on the PASCAL VOC 2007. In the first setting, we allow the randomized merging algorithm to be performed only one time in each recursion step, and the experimental results are shown in Table II. It can be seen that greedy merging and one-time randomized merging achieve comparable results. However, with the increase of randomized times, our merging algorithm generates more diverse proposals. Here we conducted the experiments and compare the results obtained with different randomized times. The number of object proposals are fixed as 500 and 1000, respectively. Figure 11 clearly shows that the AR rate improves

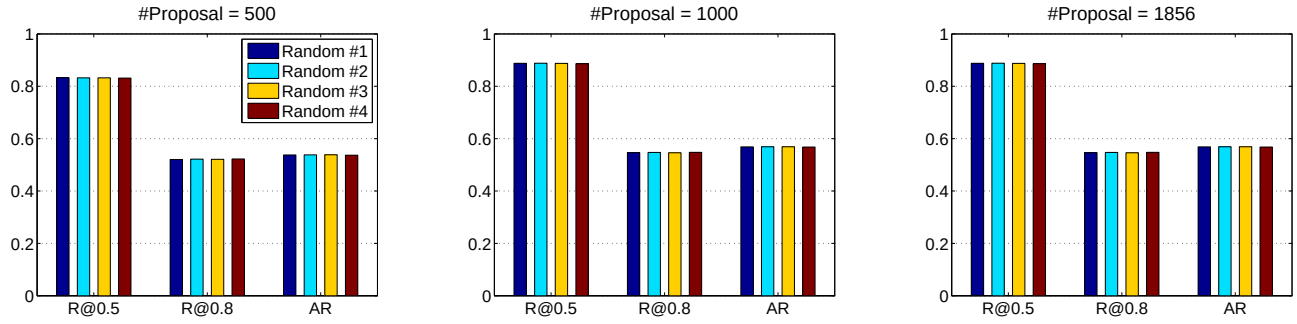

(a) (b) (c)

Fig. 12. Comparison of recall rates in four groups of randomized merging experiments on the PASCAL VOC 2007 test set. We report the results using top ranked 500, 1000 and all of the proposals. R@0.5 and R@0.8 mean the recall rates with an IoU threshold of 0.5 and 0.8, respectively, and AR is the average recall. Best viewed in color.

as the random times increases, and then goes near saturation eventually. This approach provides a notable gain in recall rates compared to greedy merging strategy. It is also shown that the recall rate with an IoU threshold of 0.5 first keep fixed and then drops as the random times increases, while that with 0.8 boosts consistently. This suggests that the predicted objectness scores are not accurate enough with small IoU values with the ground truth bounding boxes.

Another critical issue is that whether stable performance can be achieved by our proposed method, since we introduce randomness to the merging process. To better clarify this problem, we conduct four groups of experiments, and we repeat the randomized merging process for five times for each group. As shown in Figure 12, our method also exhibits great stability in recall rates and average recall with different numbers of object proposals.

VI. Conclusion

In this paper, we have presented a simple yet effective approach to hierarchically segment object proposals by develping a deep architecture of recursive neural networks. We incorporate the similarity metric learning into the bottom-up region merging process for end-to-end training, rather than manually designing various representations. In addition, we introduce randomness into the greedy search to cope with the ambiguity in the process of merging regions, making the inference more robust against noise. Extensive experiments on standard benchmarks demonstrate the superiority of our approach over state-of-the-art approaches. In addition, the effectiveness of our method for real detection systems is also verified. In future work, this proposed framework can be tightly combined with category-specific object detection methods.

References

[1] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 580–587.

[2] K. He, X. Zhang, S. Ren, and J. Sun, "Spatial pyramid pooling in deep convolutional networks for visual recognition," in *Proc. Eur. Conf. Comput. Vis.*, Sep. 2014, pp. 346–361.

[3] Y. J. Lee and K. Grauman, "Learning the easy things first: Self-paced visual category discovery," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* IEEE, Jun. 2011, pp. 1721–1728.

[4] M. Cho, S. Kwak, C. Schmid, and J. Ponce, "Unsupervised object discovery and localization in the wild: Part-based matching with bottom-up region proposals," *arXiv preprint arXiv:1501.06170*, 2015.

[5] J. R. Uijlings, K. E. van de Sande, T. Gevers, and A. W. Smeulders, "Selective search for object recognition," *Int. J. Comput. Vis.*, vol. 104, no. 2, pp. 154–171, 2013.

[6] S. Manen, M. Guillaumin, and L. Van Gool, "Prime object proposals with randomized prim's algorithm," in *Proc. IEEE Int. Conf. Comput. Vis.* IEEE, Dec. 2013, pp. 2536–2543.

[7] Y. Xiao, C. Lu, E. Tsougenis, Y. Lu, and C.-K. Tang, "Complexity-adaptive distance metric for object proposals generation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2015, pp. 778–786.

[8] R. Socher, C. C. Lin, C. Manning, and A. Y. Ng, "Parsing natural scenes and natural language with recursive neural networks," in *Proc. Int. Conf. Mach. Learn.*, Jun. 2011, pp. 129–136.

[9] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *arXiv preprint arXiv:1506.01497*, 2015.

[10] C. L. Zitnick and P. Dollár, "Edge boxes: Locating object proposals from edges," in *Proc. Eur. Conf. Comput. Vis.* Springer, Sep. 2014, pp. 391–405.

[11] M.-M. Cheng, Z. Zhang, W.-Y. Lin, and P. Torr, "BING: Binarized normed gradients for objectness estimation at 300fps," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* IEEE, Jun. 2014, pp. 3286–3293.

[12] B. Alexe, T. Deselaers, and V. Ferrari, "What is an object?" in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* IEEE, Jun. 2010, pp. 73–80.

[13] B. Alexe, T. Deselaers, and V. Ferrari, "Measuring the objectness of image windows," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 11, pp. 2189–2202, 2012.

[14] Z. Zhang, J. Warrell, and P. H. Torr, "Proposal generation for object detection using cascaded ranking svms," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011, pp. 1497–1504.

[15] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," *arXiv preprint arXiv:1411.4038*, 2014.

[16] P. Krähenbühl and V. Koltun, "Geodesic object proposals," in *Proc. Eur. Conf. Comput. Vis.* Springer, Sep. 2014, pp. 725–739.

[17] J. Carreira and C. Sminchisescu, "CPMC: Automatic object segmentation using constrained parametric min-cuts," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 7, pp. 1312–1328, 2012.

[18] P. Arbelaez, J. Pont-Tuset, J. Barron, F. Marques, and J. Malik, "Multiscale combinatorial grouping," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* IEEE, Jun. 2014, pp. 328–335.

[19] P. Rantalankila, J. Kannala, and E. Rahtu, "Generating object segmentation proposals using global and local search," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* IEEE, Jun. 2014, pp. 2417–2424.

[20] M. Bergh, G. Roig, X. Boix, S. Manen, and L. Gool, "Online video seeds for temporal window objectness," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2013, pp. 377–384.

[21] P. F. Felzenszwalb and D. P. Huttenlocher, "Efficient graph-based image segmentation," *Int. J. Comput. Vis.*, vol. 59, no. 2, pp. 167–181, 2004.

[22] R. Girshick, "Fast R-CNN," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2015, pp. 1440–1448.

[23] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.

[24] B. Taskar, D. Klein, M. Collins, D. Koller, and C. D. Manning, "Max-margin parsing." in *EMNLP*. Citeseer, 2004, p. 3.

[25] L. Bottou, "Stochastic gradient descent tricks," in *Neural Networks: Tricks of the Trade*, 2012, pp. 421–436.

[26] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *Int. J. Comput. Vis.*, vol. 88, no. 2, pp. 303–338, 2010.

[27] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, "ImageNet large scale visual recognition challenge," *Int. J. Comput. Vis.*, pp. 1–42, 2014.

[28] J. Hosang, R. Benenson, P. Dollár, and B. Schiele, "What makes for effective detection proposals?" *arXiv preprint arXiv:1502.05082*, 2015.

[29] C. Wang, L. Zhao, S. Liang, L. Zhang, J. Jia, and Y. Wei, "Object proposal by multi-branch hierarchical segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, June 2015, pp. 3873–3881.

[30] X. Chen, H. Ma, X. Wang, and Z. Zhao, "Improving object proposals with multi-thresholding straddling expansion," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2015, pp. 2587–2595.

[31] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, "Caffe: Convolutional architecture for fast feature embedding," in *Proc. ACM Multimedia*, Nov. 2014, pp. 675–678.

[32] J. Pont-Tuset, P. Arbelaez, J. Barron, F. Marques, and J. Malik, "Multiscale combinatorial grouping for image segmentation and object proposal generation." *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 1, pp. 128–140, 2016.

[33] P. Krähenbühl and V. Koltun, "Learning to propose objects," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2015, pp. 1574–1582.

[34] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, Dec. 2012, pp. 1097–1105.

Tianshui Chen received the B.E. degree from School of Information and Science Technology, Sun Yat-sen University, Guangzhou, China, in 2013, where he is currently pursuing the Ph.D. degree in computer science with the School of Data and Computer Science. His current research interests include computer vision and machine learning.

Liang Lin (M'09, SM'15) is a full Professor with the School of Data and Computer science, Sun Yat-Sen University (SYSU), China. He received the B.S. and Ph.D. degrees from the Beijing Institute of Technology (BIT), Beijing, China, in 1999 and 2008, respectively. From 2008 to 2009, he was a jPost-Doctoral Research Fellow with the Department of Statistics, University of California, Los Angeles (UCLA). His research focuses on new models, algorithms and systems for intelligent processing and understanding of visual data such as images and videos. He has authorized and co-authorized on more than 100 papers in top tier academic journals and conferences. He has been serving as an associate editor of IEEE Trans. Human-Machine Systems. He was supported by several promotive programs or funds for his works such as the NSFC Excellent Young Scholars Program in 2016. He was the recipient of the Best Paper Runners-Up Award in ACM NPAR 2010, Google Faculty Award in 2012, and Best Student Paper Award in IEEE ICME 2014.

Xian Wu received the B.E. degree from the School of Software, Sun Yat-sen University, Guangzhou, China, in 2015, where he is currently pursuing the MS c. degree in Software Engineering with the School of Data and Computer Science. His current research interests include computer vision and machine learning.

Xiaonan Luo is a professor of the School of Information and Science Technology at Sun Yat-sen University. He also serves as Director of the National Engineering Research Center of Digital Life. His research interests include Computer Vision, Image Processing, and Computer Graphics & CAD.